

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping

INTRODUCTION

The Optical Transport Network (OTN) emerged in the late 1990s as a “digital wrapper” around client signals before transport over a WDM network. Its main goals were to provide a common way of supervising different client signals, deliver reliability comparable to or better than the existing SONET/SDH networks, and provide the ability to build a carrier’s carrier network.

Standardization efforts by the ITU-T led to the approval of a first version in 2001 optimized for the SONET/SDH clients¹ which were dominant at that time. Though this standard did provide connectivity at the sub-wavelength level, through the optical channel data unit (ODU) container, it was believed that most networking would occur at the wavelength level.

Since then, the telecommunication world has changed dramatically. Migration towards packet-based networks introduced Ethernet as natural client signals for the OTN. The ITU-T faced the challenge of efficiently supporting these new client signals while retaining backwards compatibility with the installed base. Proposed solutions introduced lower order containers, creating the need for networking at the sub-wavelength (i.e. electrical) ODU level.

This paper discusses these lower order ODU0 and ODUflex containers and how they fit into the OTN hierarchy.

EFFICIENT TRANSPORT OF GbE: ODU0

The rapidly growing number of GbE interfaces in metro networks increases the need for a cost effective transparent transport of GbE over OTN. Transport equipment already supports GbE over SONET/SDH virtual concatenated containers. However, this solution would not allow carriers to remove SONET/SDH from their networks. Therefore an OTN specific solution was needed.




Figure 1: GbE Mapped into ODU1

Though existing standards already allowed mapping of GbE directly into ODU1 using transparent GFP (GFP-T), this would create a waste of bandwidth, see Figure 1. This

situation was not deemed acceptable for such a popular client signal. Multiplexing two GbE signals into the same ODU1, using GFP’s linear extension mode, results in the disadvantage of losing client-level supervision and only allows networking at the undesirable GFP level.




Figure 2: ODU0 Increases Bandwidth Efficiency

Therefore, a new ODU0 container was introduced² which exactly fit half the payload area of an ODU1 container. Since each ODU0 can transparently map one GbE client, the capacity for GbE transport is effectively doubled. By definition, this ODU0 is a lower-order ODUk for sub-wavelength networking – an OTU0 does not exist.

Frame Format

One of the properties of OTN is that while the frame rate changes over the different levels in the OTN hierarchy, the frame format remains identical. The ODU0 frame format is no exception. Like all other ODUks it consists of a structure of four rows and 3824 columns, as presented in Figure 3.




Figure 3: ODUk Frame Format

This structure is further divided into the ODUk overhead area (the first fourteen columns) and the optical channel payload unit (OPU) area. The latter contains two columns of overhead and 3808 columns of payload area which is available for the mapping of client data. Evidently, the OTUk overhead area in the ODU0 frame will remain unused.

The nominal ODU0 rate equals half the payload area rate of an ODU1. The latter is tailored for transport of STM-16/OC-48 signals at 2,488.32 Mbit/s. From this we can derive the ODU0 rate to be 1,244.16 Mbit/s ± 20 ppm, while the rate of the available OPU0 payload area is 1,238.95431 Mbit/s.

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping

Payload Compression

The line rate of an 8B/10B encoded GbE signal is 1.25 Gbit/s, slightly higher than the available OPU0 payload area. Fortunately, the signal can be compressed: GFP-T mapping already defines transcoding of 8B/10B to a 64B/65B code³. If, in this process, no rate adaptation (idle frames or padding characters) is performed, the resulting stream of GFP-T frames will still be synchronous to the original signal. The parameters chosen for this timing transparent transcoding (TTT) compress the GbE signal by a factor of 15/16.

GbE Client Mapping

An old idea was revived to address the mapping of the compressed GbE client signal into the OPU0 payload area. The generic mapping procedure (GMP) was first proposed during the initial standardization phase of OTN, but was dismissed at the time because it lacked applicable client signals. GMP allows mapping of any constant bit rate (CBR) client into the OPUk payload area by evenly spreading the client bytes over the available OPUk payload positions. In between, stuffing bytes are introduced to provide positive justification. Their locations are not predefined, but are calculated in real-time based on the actual incoming number of client bits, thus minimizing mapping jitter.

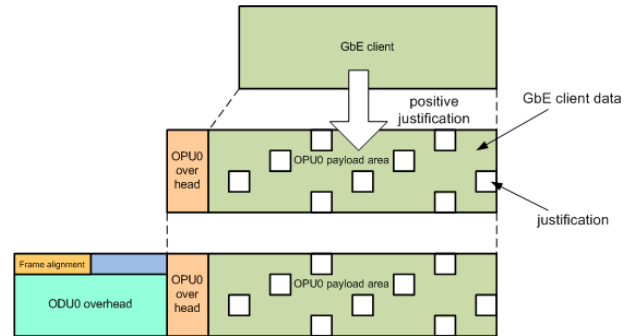


Figure 4: Generic Mapping Procedure (GMP)

Since GMP can only perform positive justification, the upper limit of the CBR client rate has to be lower than the lower limit of the rate of available payload bits in the server signal - in this case the OPUk payload area.

ODU0 Mapping

In order to multiplex ODU0s into an ODU1, the payload area of the latter has been divided into two time slots called optical channel tributary unit 0 into 1 (ODTU01). As shown in Figure 5, each ODU0 is mapped into an ODTU01 time slot using the asynchronous mapping procedure (AMP), which is consistent with the legacy mapping of ODUj into ODUk.

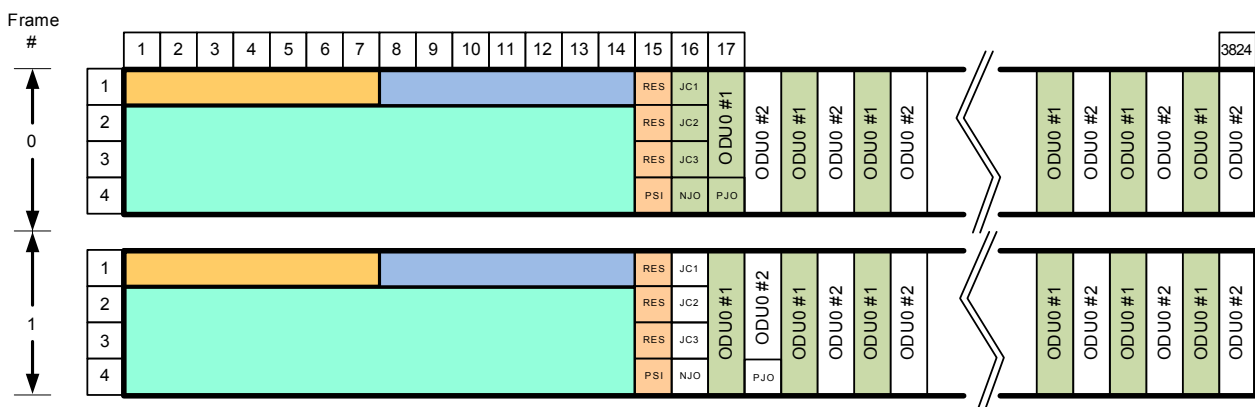


Figure 5: Mapping of ODU0 into OPU1 Time Slots

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping

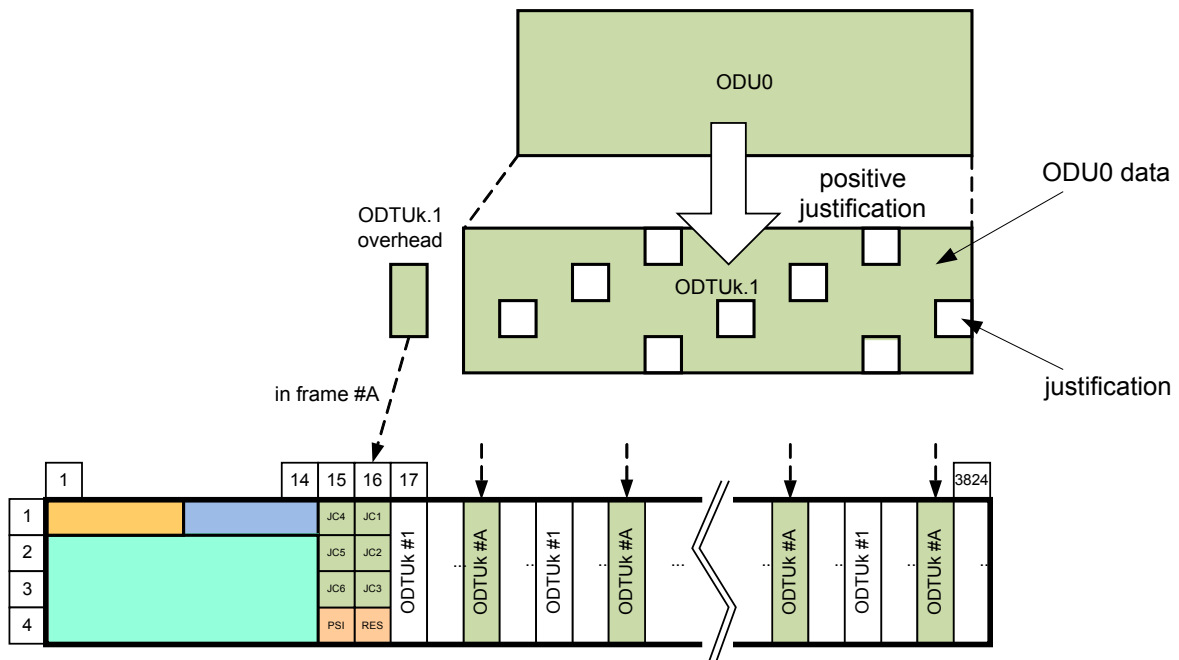


Figure 6: Mapping of ODU0 into an OPUk Time Slot ($k < 1$)

Legacy higher order ODUk were substructured into time slots of approximately 2.5 Gbit/s, nicely fitting the OTN wrappers around SONET/SDH signals. More recently, the ODUk levels above ODU1 are also subdivided into time slots at the approximate rate of 1.25 Gbit/s to support multiplexing of ODU0s into a single stage⁴. These time slots are simply called ODTUk since they have more uses beyond mapping of ODU0s. However, AMP only works for mapping of clients with a rate very close to the ODTU rate. Rather than defining fixed stuff columns, the ITU-T opted to use the GMP mechanism to map ODU0s into a single ODTU time slot. Figure 6 shows an ODU0 allocated to ODTU time slot number A. The associated GMP overhead is inserted into the OPUk overhead once every multiframe.

GbE Mapping

Figure 7 summarizes the steps needed to map GbE signals into a higher-order ODUk. First, the 8B/10B encoded GbE signal has to be compressed to fit into the space made available by an ODU0. This compressed GbE is then mapped into the payload area of an OPU0 and then ODUk overhead is added. The resulting ODU0 is then AMP mapped into an OPU1 time slot or GMP mapped into an OPUk ($k > 1$) time slot before being multiplexed together with other time slots into the payload area of the higher order OPUk. Finally, the higher-order ODUk overhead is added.

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping

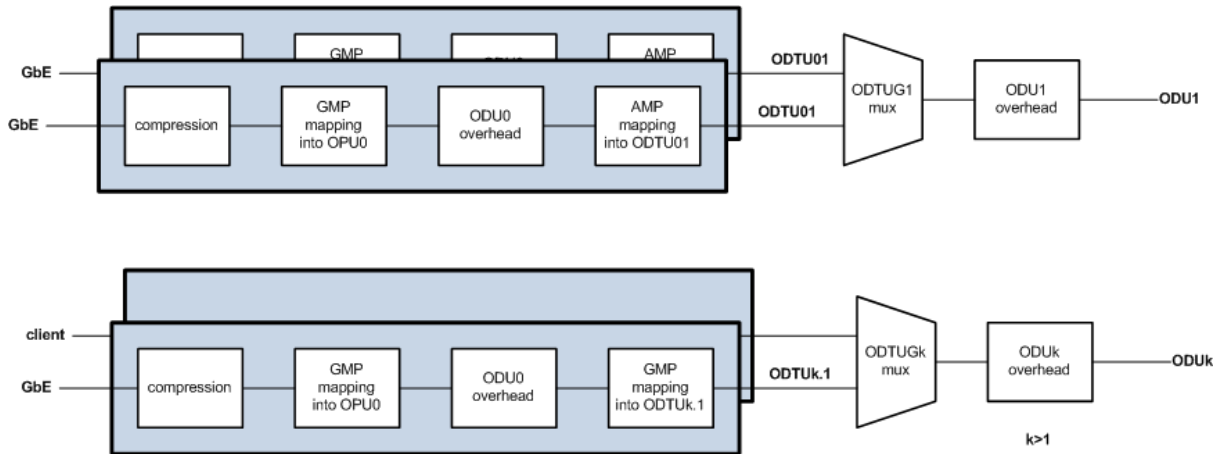


Figure 7: GbE Mapping

ODU0 Compared to VC-12 / VT1.5 SPE

Since an ODU0 has no associated OTUk signal, it is, by definition, a lower-order ODUk that needs to be multiplexed into a higher-order ODUk. In order to support networking at the ODU0 level, future OTN network elements must be able to provide ODU0 add/drop and protection features.

The resulting network architecture bears a strong resemblance to the SONET/SDH networks with VC-12/VT1.5 SPE lower-order signals multiplexed into VC-4/STS-1 SPE higher-order signals.

Other Sub-ODU1 Clients

The techniques introduced to optimize the transport of GbE over the OTN have been reused⁴. GMP allows mapping of clients with a rate above 1.238 Gbit/s but below 2.488 Gbit/s into ODU1 independent of the client rate. An example of such a client is FC-200 Fibre Channel (2G FC). Furthermore, client signals with a rate below 1.238 Gbit/s, like STM-1/OC-3, STM-4/OC-12, and FC-100 (1G FC), can be mapped into an ODU0 through GMP. These signals, which are not expected to become sufficiently important as clients for the OTN, would not typically justify further standardized optimization. In those few applications where it does make sense, proprietary solutions already exist. For example, TPACK's TPO124 and TPO114 OTN mapper devices allow efficient point-to-point transport of sub-ODU1 client signals by further dividing the OPU1 payload area into sixteen 155.52 Mbit/s time slots.

FLEXIBLE CONTAINER FOR ALL: ODUFLEX

The experience with Ethernet clients showed that it is difficult to predict the rate of future client signals. In the wake of Ethernet a number of other candidate clients appeared, such as Fibre Channel and video distribution signals. Most of these would not fit into any existing ODUk without significant loss of bandwidth. Also, defining a new ODU container each time a new client is added was considered impractical. Moreover, it would require upgrading ODUk switch fabrics. Using ODUk virtual concatenation to map CBR clients was viewed as being too complex and does not provide unique mappings; e.g. a 12 Gbit/s client could be mapped into either an ODU1-5v or an ODU2-2v container.

The solution was the introduction of a flexible lower order ODUflex container that can be "right sized" to fit any client rate and thus occupies the minimum number of time slots in the higher order ODUk.

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping




Figure 8: Mapping a FC-400 (4GFC) into Four OPU2 Time Slots

As an example, Figure 8 shows how an FC-400 Fibre Channel (4G FC) is mapped into an ODUflex occupying four OPU2 time slots. This leaves the other four OPU2 time slots available for other use.

Frame Format

The ODUflex frame is consistent with the generic ODUk frame depicted in Figure 3.

Client Mapping

CBR clients are mapped into the ODUflex by using the bit-synchronous mapping procedure (BMP). This procedure is very simple. As shown in Figure 9, it simply involves wrapping ODUk/OPUk overhead around the client data. Thus, the latter exactly fits the OPUflex payload area without the need for justification.

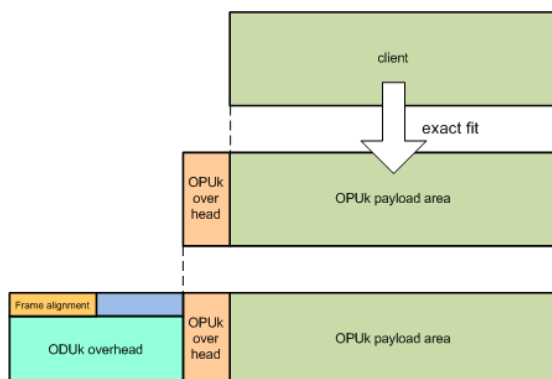


Figure 9: Bit-Synchronous Mapping Procedure (BMP)

The ODUflex is therefore the exact right size for its client: its bit rate depends directly on the client bit rate

$$f_{ODUflex} = \frac{239}{238} \times f_{client}$$

The ODUflex can also be used as a flexible rate container for GFP-F mapped packets. Though the container rate can be chosen arbitrarily, there is no real advantage in selecting rates that do not correspond to a multiple of higher order OPUk time slots.

ODUflex Mapping

To map an ODUflex into a higher-order ODUk, the appropriate number *ts* of higher-order OPUk time slots is allocated. The selection of time slots is arbitrary, though all time slots should be part of the same ODUk. The resulting structure is called an ODTUK.*ts*. Since both the ODUflex rate and the ODTUK.*ts* rate can vary, GMP is the natural choice for the mapping procedure.

Figure 10 shows an ODUflex allocated to four time slots numbered A, B, C and D. The GMP overhead is associated with the last time slot and is inserted into the OPUk overhead once every multiframe.

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping

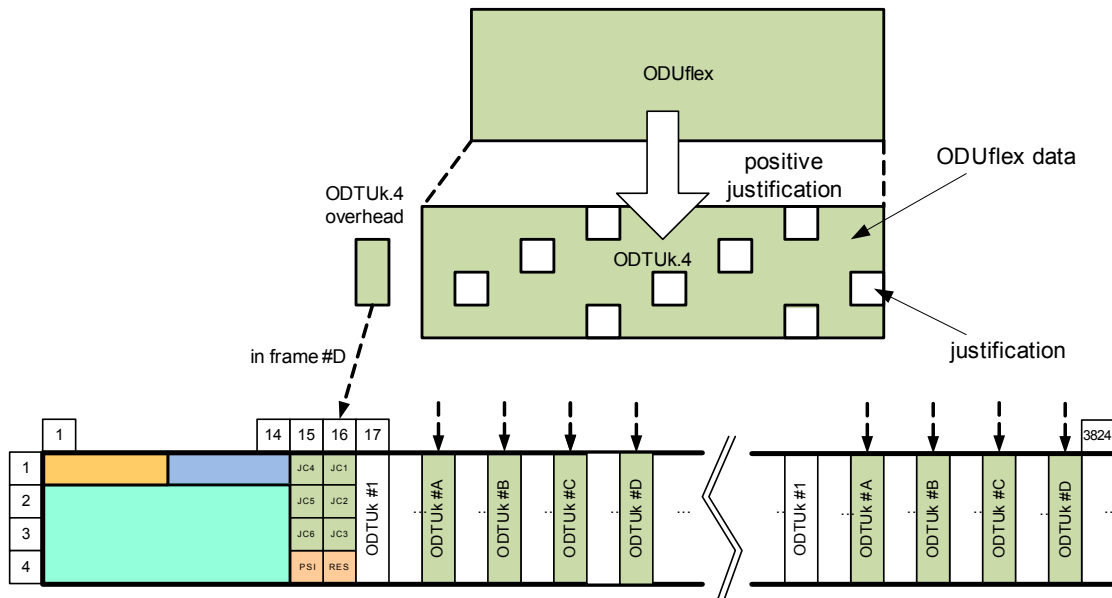


Figure 10: Mapping of ODUflex into Four OPUk Time Slots

4G FC Mapping

Figure 11 shows an example of mapping a FC-400 (4G FC) signal into an ODU2 through an ODUflex. First, the client signal is bit-synchronously mapped into the OPUflex payload area after which ODUk overhead is added. The resulting ODUflex is, in turn, GMP mapped into a structure of four ODTU2 time slots or ODTU2.4. These four time slots are multiplexed together into the OPU2 payload area with other time slots carrying other ODTUs and finally ODU2 overhead is added.

ODUflex versus Concatenation

ODUflex shares properties with both contiguous and virtual concatenation techniques used in SONET/SDH networks.

Figure 12 shows how a client signal is rate-adapted into a contiguous concatenation group (CCG) container that can have a discrete number of pre-defined sizes, all exact multiples of the time slot size. The CCG is then mapped into adjacent time slots and routed as a single signal through the network requiring support at each intermediate node.

The size of a virtual concatenation group (VCG) can be any multiple of the time slot size. The client is rate-adapted into the VCG container. The latter is then inverse multiplexed into arbitrary time slots which can be independently routed through the network and only require support at the access nodes.




Figure 11: 4G FC Mapping

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping

With an ODUflex the client exactly fits the payload container. Therefore, its size depends solely on the client and is independent of the time slot size. As with contiguous concatenation, the time slots allocated for an ODUflex need to be a part of the same physical signal. Like virtual concatenation, the time slots can be arbitrarily allocated. The ODUflex is rate-adapted into the time slots and routed as a single signal through the network. Since it is a lower order ODU it can be tunneled through legacy OTN nodes.

When to Use ODUflex?

It would be tempting to consider ODUflex as the generic solution for all future CBR client mapping. However, this does not always make sense because client signals with a rate close enough to an existing fixed rate ODUk can be more efficiently mapped directly into this ODUk. The ODUflex is only preferred if it will occupy fewer time slots than the smallest non-flexible ODUk capable of fitting the client signal.




Figure 12: ODUflex versus Concatenation

Table 1 provides an overview of the ODUflex properties compared to contiguous and virtual concatenation.

Table 1: ODUflex properties versus concatenation properties

	ODUflex	CCG SONET/SDH	VCG	
			SONET/SDH	OTN
# of path signals	1	1	X (inverse multiplexed)	
Bandwidth	continuous	multiple of time slot bandwidth	multiple of time slot bandwidth	
CBR clients	yes	GFP-T	GFP-T	
# time slots	1.. max # of HO ODUk time slots	4, 16, 64, 256 HO time slots	1..64 (LO) 1..256 (HO)	1..256
Time slot selection	arbitrary within the same HO ODUk	adjacent	arbitrary	
Diverse routing	no	no	yes	
Delay compensation needed	no	no	yes	

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping



For example, an ODUflex container for an FC-800 (8G FC) client signal fits into 7 ODU2 time slots, leaving the 8th time slot available for an additional client. In contrast, even though an ODUflex wrapper around a 100GbE signal would fit into an ODU4, it would occupy all 80 time slots and therefore does not make sense. Instead, 100GbE is directly mapped into ODU4 through GMP.

Another ODUflex application is the transport of GFP-F encapsulated packet streams like Ethernet, MPLS-TP or IP/MPLS, providing resizeable tunnels between L2/L3 switch fabrics. These ODUflexes make most efficient use of the OTN bandwidth when configured in exact multiples of the lowest rate ODTUk.1 in the network.

CONCLUSION

To address the growing number of new client signals, ITU-T added two new lower order ODU containers to the OTN hierarchy. ODU0 is optimized for GbE clients, but also allows transport of client signals with a rate below 1.238 Gbit/s. ODUflex provides a future-proof solution for any client with a rate above ODU1. A generic mapping procedure (GMP) was developed which is suitable both for mapping of clients into fixed sized ODUs and for mapping of lower-order ODUs into higher-order ODU time slots.

Though both ODU0 and ODUflex are defined consistent with existing ODUk signals, their introduction creates the need for networking at the ODUk level.

These developments add important requirements to next-generation OTN equipment: support for ODU0 and ODUflex, GMP capable multiplexing of lower order ODUs into 1.25 Gbit/s higher-order ODU time slots and lower-order ODU cross-connecting.

The techniques introduced are being reused:

- From the ODU4 level onwards all lower order ODUs are mapped using GMP into ODTU4 time slots.
- The ODU2e container introduced as an extension to optimize 10GbE LAN transport can be pulled into the standard hierarchy by treating it similar to ODUflex.
- Sub-ODU1 clients can be mapped into ODU0 or ODU1 using GMP.

ODU0 and ODUflex

A Future-Proof Solution for OTN Client Mapping



ABBREVIATIONS

Abbrev	Description
AMP	Asynchronous mapping procedure
BMP	Bit-synchronous mapping procedure
CBR	Constant Bit Rate
CCG	Contiguous Concatenation Group
FC	Fibre Channel
GbE	Gigabit Ethernet
GFP-T	Transparent Generic Framing Procedure
GMP	Generic mapping procedure
HO	Higher order
ITU-T	International Telecommunication Union – Standardization Sector
LO	Lower order
OC	Optical Carrier
ODU	Optical channel Data Unit
ODTUG	Optical channel Data Tributary Unit Group
ODTUjk	Optical channel Data Tributary Unit j into k
ODTUK.ts	Optical channel Data Tributary Unit k with ts tributary slots
OPU	Optical channel Payload Unit
OTN	Optical Transport Network
OTU	Optical channel Transport Unit
TTT	Timing Transparent Transcoding
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
SPE	Synchronous Payload Envelope
STM	Synchronous Transport Module
VC	Virtual Container
VCG	Virtual Concatenation Group
VT	Virtual Tributary
WDM	Wavelength Division Multiplexing

REFERENCES

- ¹ ITU-T Recommendation G.709/Y.1331 (02/01), Interfaces for the optical transport network (OTN), ITU-T, Geneva, February, 2001.
- ² ITU-T Recommendation G.709/Y.1331 (2003) Amendment 3 (04/09), Interfaces for the optical transport network (OTN), ITU-T, Geneva, April, 2009.
- ³ ITU-T Recommendation G.7041/Y.1303 (10/08), Generic framing procedure (GFP), ITU-T, Geneva, October, 2008.
- ⁴ ITU-T Recommendation G.709/Y.1331 (12/09), Interfaces for the optical transport network (OTN), ITU-T, Geneva, December, 2009.